Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	0	((instance near2 data) or datagate) same metadata same metaview	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/26 16:41
L2	20207	((instance near2 data) or datagate)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/26 16:41
L3	1	2 and metadata and metaview	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/26 17:54
L4	428	2 and metadata and directory	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/26 17:54
L5	102	4 and dictionary	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/26 17:55
L6	54	5 and (xml with document)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/26 17:55
L7	49	6 and authentication	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/26 18:43
L8	23498	"707"/$.ccls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/26 17:56
L9	32392	"709"/$.ccls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/26 17:56
L10	3103	8 and 9	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/26 17:58

L11	7	7 and 10	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/26 18:02
L12	2	(adm or (adaptive near data near management)) same (metadata or (meta near data) same model)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/26 18:40
L13	43	(adm or (adaptive near data near management)) and (metadata or (meta near data) same model)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/26 18:40
L14	40	13 and (tree or graph)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/26 18:45
L15	50311	access$3 same security	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/26 18:42
L16	4914	user$1 same (access near right$1)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/26 18:43
L17	2536	15 and 16	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/26 18:43
L18	1572	17 and authenticat$4	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/26 18:44
L19	15469	relational near database	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/26 18:45
L20	322	18 and 19	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/26 18:45

L21	259394	(relationship or relation$5) and (tree or graph)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/26 18:45
L22	259602	20 ansd 21	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/26 18:46
L23	210	20 and 21	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/26 18:46
L24	112	20 not 21	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/26 18:46
L25	3	10 and 24	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2004/11/26 18:46

	Document ID	Kind Codes	Source	Issue Date	Pages
1	US 6701345 B1		USPAT	20040302	96
2	US 6584466 B1		USPAT	20030624	26
3	US 6314425 B1		USPAT	20011106	28

	Title	Abstract
1	Providing a notification when a plurality of users are altering similar data in a health care solution environment	
2	Internet document management system and methods	
3	Apparatus and methods for use of access tokens in an internet document management system	

	Current OR	Current XRef	Retrieval Classif	Inventor
1	709/205	707/8; 709/232		Carley; Kevin W. et al.
2	707/10	707/203; 707/3; 707/5; 709/226; 715/501.1		Serbinis; M. Michael et al.
3	707/10	705/26; 705/8; 707/104.1; 709/218; 709/224; 711/121; 713/200; 715/501.1; 715/513		Serbinis; M. Michael et al.

◆IEEE

Membership   Publications/Services   Standards   Conferences   Careers/Jobs

IEEE Xplore® RELEASE 1.8

Welcome
United States Patent and Trademark Office

IEEE Xplore®
1 Million Documents.
1 Million Users
...And Growing

Help   FAQ   Terms   IEEE Peer Review   **Quick Links**   ▣

» Search Results

Welcome to IEEE Xplore®

○ Home
○ What Can I Access?
○ Log-out

Tables of Contents

○ Journals & Magazines
○ Conference Proceedings
○ Standards

Search

○ By Author
○ Basic
○ Advanced
○ CrossRef

Member Services

○ Join IEEE
○ Establish IEEE Web Account
○ Access the IEEE Member Digital Library

IEEE Enterprise

○ Access the IEEE Enterprise File Cabinet

Your search matched **2** of **1097671** documents.
A maximum of **500** results are displayed, **15** to a page, sorted by **Relevance** in **Descending** order.

Refine This Search:
You may refine your search by editing the current search expression or entering a new one in the text box.

| authenticat* and (access* <near/3> right*) and datal | **Search** |

☐ Check to search within this result set

Results Key:
JNL = Journal or Magazine   **CNF** = Conference   **STD** = Standard

---

1 **Authentication mechanisms in microprocessor-based local area networks**
Ciminiera, L.; Valenzano, A.;
Software Engineering, IEEE Transactions on , Volume: 15 , Issue: 5 , May 1989
Pages:654 - 658

[Abstract]   [PDF Full-Text (520 KB)]   IEEE JNL

---

2 **Distributed file systems: focus on Andrew File System/Distributed File Service (AFS/DFS)**
Tobbicke, R.;
Mass Storage Systems, 1994. 'Towards Distributed Storage and Data Management Systems.' First International Symposium. Proceedings., Thirteenth IEEE Symposium on , 12-16 June 1994
Pages:23 - 26

[Abstract]   [PDF Full-Text (296 KB)]   IEEE CNF

🖨 **Print Format**

IEEE HOME I SEARCH IEEE I SHOP I WEB ACCOUNT I CONTACT IEEE

◆IEEE

Membership   Publications/Services   Standards   Conferences   Careers/Jobs

IEEE *Xplore*®
RELEASE 1.8

Welcome
United States Patent and Trademark Office

IEEE *Xplore*®
1 Million Documents
1 Million Users
...And Growing

» ABSTRACT PLUS

Help   FAQ   Terms   IEEE Peer Review

Quick Links

Welcome to IEEE *Xplore*

○ Home
○ What Can
    I Access?
○ Log-out

Tables of Contents

○ Journals
    & Magazines
○ Conference
    Proceedings
○ Standards

Search

○ By Author
○ Basic
○ Advanced
○ CrossRef

Member Services

○ Join IEEE
○ Establish IEEE
    Web Account
○ Access the
    IEEE Member
    Digital Library

IEEE Enterprise

○ Access the
    IEEE Enterprise
    File Cabinet

🖨 Print Format

Search Results   [PDF FULL-TEXT 520 KB]   NEXT   DOWNLOAD CITATION

Request Permissions
RIGHTSLINK⟨⟩
Copyright Clearance Center, Inc.

Authentication mechanisms in microprocessor-based local area networks

Ciminiera, L.   Valenzano, A.
Dipartimento di Autom. e Inf., Politecnico di Torino, Italy;
This paper appears in: **Software Engineering, IEEE Transactions on**

Abstract:
The problem of **authenticating** the users of a computer network in order to protect the shared resources against unauthorized use is discussed. Since intruders could enter the network and try to use services they have no **right** to **access**, the host implementing the service (or server) has to check the user's identity and **access rights** by searching in the relevant **database**. The author presents a method of carrying out such checks efficiently. The basic idea is that a suitable interface process is associated with each user-server connection in order to filter out unauthorized requests, thus implementing a sort of cache with parallel search where the working set of the whole **database** is stored and explored. The use of the interface process enables the system to exploit the hardware support for capability checking provided by new microprocessors. In particular, an implementation using iAPX432-based hosts is illustrated and performance issues are discussed

Index Terms:
local area networks   security of data   **access rights**   **authentication** mechanisms   cache
capability checking   check   **database**   hardware support   iAPX432-based hosts   identity   interface
process   intruders   microprocessor-based local area networks   parallel search   performance
issues   protect   searching   server   service   shared resources   unauthorized requests
unauthorized use   user-server connection

Documents that cite this document
There are no citing documents available in IEEE Xplore at this time.

Search Results   [PDF FULL-TEXT 520 KB]   NEXT   DOWNLOAD CITATION

Search Results   [PDF FULL-TEXT 296 KB]   PREV   DOWNLOAD CITATION

Request Permissions
RIGHTSLINK⟨⟩

Distributed file systems: focus on Andrew File System/Distributed File Service (AFS/DFS)

Tobbicke, R.
CERN, Geneva, Switzerland;

This paper appears in: **Mass Storage Systems, 1994. 'Towards Distributed Storage and Data Management Systems.' First International Symposium. Proceedings., Thirteenth IEEE Symposium on**
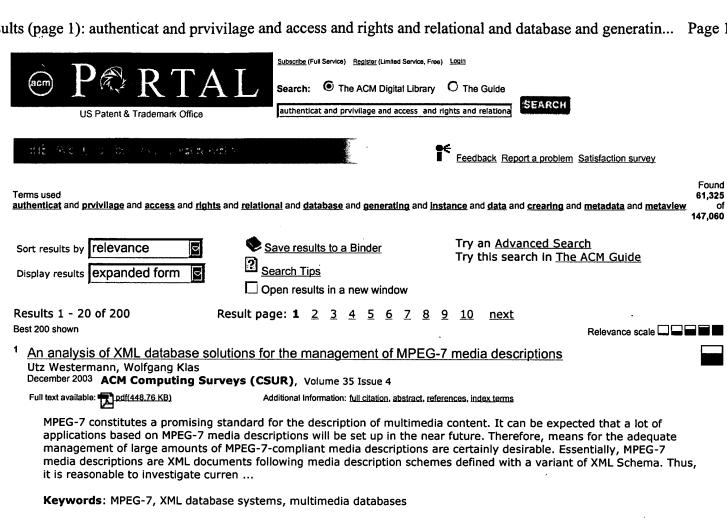
Abstract:
Currently, there is a move away from huge, monolithic mainframe computers to smaller, individual workstations. This move requires that data and programs be shared between individual users on the same computer as well as across machines on a local area network (LAN) or even on a wide area network (WAN). A popular distributed file system is Transarc's Andrew File System (AFS), the precursor of the Open Software Foundation's (OSF) Distributed File Service (DFS), part of OSF's Distributed Computing Environment. This paper compares AFS and DFS with other successful distributed file systems, discussing clients and file servers, data cache (cache consistency and **authentication**), file system topology, **access rights**, privileged programs, critical data, and **authentication** in long-running batch jobs

Index Terms:
AFS/DFS   Andrew File System/Distributed File Service   Distributed Computing Environment   OSF Open Software Foundation   Transarc   **access rights**   **authentication**   buffer storage   cache consistency   clients   critical data   data cache   distributed **databases**   distributed file system   file servers   file system topology   local area network   long-running batch jobs   message **authentication**   privileged programs   security of data   wide area network   workstations
AFS/DFS   Andrew File System/Distributed File Service   Distributed Computing Environment   OSF Open Software Foundation   Transarc   **access rights**   **authentication**   buffer storage   cache consistency   clients   critical data   data cache   distributed **databases**   distributed file system   file servers   file system topology   local area network   long-running batch jobs   message **authentication**   privileged programs   security of data   wide area network   workstations

Documents that cite this document

Terms used	Found 61,325
authenticat and **prvivilage** and **access** and **rights** and **relational** and **database** and **generating** and **instance** and **data** and **crearing** and **metadata** and **metaview**	of 147,060

Sort results by [relevance ▼]

Display results [expanded form ▼]

❧ Save results to a Binder

🔲 Search Tips

☐ Open results in a new window

Try an Advanced Search
Try this search in The ACM Guide

Results 1 - 20 of 200
Best 200 shown

Result page: **1** 2 3 4 5 6 7 8 9 10 next

Relevance scale ☐▢▬▰■

1 An analysis of XML database solutions for the management of MPEG-7 media descriptions ■

Utz Westermann, Wolfgang Klas

December 2003 **ACM Computing Surveys (CSUR)**, Volume 35 Issue 4

Full text available: 📄 pdf(448.76 KB)   Additional Information: full citation, abstract, references, index terms

MPEG-7 constitutes a promising standard for the description of multimedia content. It can be expected that a lot of applications based on MPEG-7 media descriptions will be set up in the near future. Therefore, means for the adequate management of large amounts of MPEG-7-compliant media descriptions are certainly desirable. Essentially, MPEG-7 media descriptions are XML documents following media description schemes defined with a variant of XML Schema. Thus, it is reasonable to investigate curren ...

Keywords: MPEG-7, XML database systems, multimedia databases

2 Semantic integrity support in SQL:1999 and commercial (object-)relational database management systems ■

Can Türker, Michael Gertz

December 2001 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 10 Issue 4

Full text available: 📄 pdf(345.55 KB)   Additional Information: full citation, abstract, index terms

The correctness of the data managed by database systems is vital to any application that utilizes data for business, research, and decision-making purposes. To guard databases against erroneous data not reflecting real-world data or business rules, semantic integrity constraints can be specified during database design. Current commercial database management systems provide various means to implement mechanisms to enforce semantic integrity constraints at database run-time.In this paper, we give ...

Keywords: Constraint enforcement, Object-relational databases, SQL:1999, Semantic integrity constraints

3 Fast detection of communication patterns in distributed executions ■

Thomas Kunz, Michiel F. H. Seuren

November 1997 **Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research**

Full text available: 📄 pdf(4.21 MB)   Additional Information: full citation, abstract, references, index terms

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial commun ...

4 Query evaluation techniques for large databases ■

Goetz Graefe

June 1993 **ACM Computing Surveys (CSUR)**, Volume 25 Issue 2

Full text available: 📄 pdf(9.37 MB)   Additional Information: full citation, abstract, references, citings, index terms, review

Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. The advent of object-oriented and extensible database systems will not solve this problem. On the contrary, modern data models exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's database systems manipulate

simple records, query-processi ...

Keywords: complex query evaluation plans, dynamic query evaluation plans, extensible database systems, iterators, object-oriented database systems, operator model of parallelization, parallel algorithms, relational database systems, set-matching algorithms, sort-hash duality

⁵ <u>Query processing over object views of relational data</u>
Gustav Fahl, Tore Risch
November 1997   **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 6 Issue 4
Full text available: 📄 <u>pdf(454.31 KB)</u>   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>index terms</u>

This paper presents an approach to *object view* management for relational databases. Such a view mechanism makes it possible for users to transparently work with data in a relational database as if it was stored in an object-oriented (OO) database. A query against the object view is translated to one or several queries against the relational database. The results of these queries are then processed to form an answer to the initial query. The approach is not restricted to a 'pure&rsqu ...

Keywords: Object views, Object-oriented federated databases, Query optimization, Relational databases, query processing

⁶ <u>The model-assisted global query system for multiple databases in distributed enterprises</u>
Waiman Cheung, Cheng Hsu
October 1996   **ACM Transactions on Information Systems (TOIS)**, Volume 14 Issue 4
Full text available: 📄 <u>pdf(697.73 KB)</u>   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index terms</u>

Today's enterprises typically employ multiple information systems, which are independently developed, locally administered, and different in logical or physical designs. Therefore, a fundamental challenge in enterprise information management is the sharing of information for enterprise users across organizational boundaries; this requires a global query system capable of providing on-line intelligent assistance to users. Conventional technologies, such as schema-based query languages and ha ...

⁷ <u>Contributed articles: Resource description framework: metadata and its applications</u>
K. Selçuk Candan, Huan Liu, Reshma Suvarna
July 2001   **ACM SIGKDD Explorations Newsletter**, Volume 3 Issue 1
Full text available: 📄 <u>pdf(1.02 MB)</u>   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>

Universality, the property of the Web that makes it the largest data and information source in the world, is also the property behind the lack of a uniform organization scheme that would allow easy access to data and information. A semantic web, wherein different applications and Web sites can exchange information and hence exploit Web data and information to their full potential, requires the information about Web resources to be represented in a detailed and structured manner. Resource Descrip ...

Keywords: Resource Description Framework (RDF), Web, XML, metadata, semantic web

⁸ <u>Model-driven development of Web applications: the AutoWeb system</u>
Piero Fraternali, Paolo Paolini
October 2000   **ACM Transactions on Information Systems (TOIS)**, Volume 18 Issue 4
Full text available: 📄 <u>pdf(6.94 MB)</u>   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index terms</u>

This paper describes a methodology for the development of WWW applications and a tool environment specifically tailored for the methodology. The methodology and the development environment are based upon models and techniques already used in the hypermedia, information systems, and software engineering fields, adapted and blended in an original mix. The foundation of the proposal is the conceptual design of WWW applications, using HDM-lite, a notation for the specification of structure, nav ...

Keywords: HTML, WWW, application, development, intranet, modeling

⁹ <u>Structured document storage and refined declarative and navigational access mechanisms in HyperStorM</u>
Klemens Böhm, Karl Aberer, Erich J. Neuhold, Xiaoya Yang
November 1997   **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 6 Issue 4
Full text available: 📄 <u>pdf(184.18 KB)</u>   Additional Information: <u>full citation</u>, <u>abstract</u>, <u>index terms</u>

The combination of SGML and database technology allows to refine both declarative and navigational access mechanisms for structured document collection: with regard to declarative access, the user can formulate complex information needs without knowing a query language, the respective document type definition (DTD) or the underlying modelling.

Navigational access is eased by hyperlink-rendition mechanisms going beyond plain link-integrity checking. With our approach, the database-internal repres ...

Keywords: Document query languages, Navigation, OODBMSs, SGML

10 Modeling the storage architectures of commercial database systems
D. S. Batory
December 1985 **ACM Transactions on Database Systems (TODS)**, Volume 10 Issue 4

Full text available: pdf(4.46 MB)   Additional Information: full citation, abstract, references, citings, index terms, review

Modeling the storage structures of a DBMS is a prerequisite to understanding and optimizing database performance. Previously, such modeling was very difficult because the fundamental role of conceptual-to-internal mappings in DBMS implementations went unrecognized. In this paper we present a model of physical databases, called the transformation model, that makes conceptual-to-internal mappings explicit. By exposing such mappings, we show that it is possible to model the storage ...

11 Early user---system interaction for database selection in massive domain-specific online environments
Jack G. Conrad, Joanne R. S. Claussen
January 2003 **ACM Transactions on Information Systems (TOIS)**, Volume 21 Issue 1

Full text available: pdf(845.54 KB)   Additional Information: full citation, abstract, references, index terms

The continued growth of very large data environments such as Westlaw and Dialog, in addition to the World Wide Web, increases the importance of effective and efficient database selection and searching. Current research focuses largely on completely autonomous and automatic selection, searching, and results merging in distributed environments. This fully automatic approach has significant deficiencies, including reliance upon thresholds below which databases with relevant documents are not search ...

Keywords: Database selection, metadata for retrieval, structuring information to aid search and navigation, user interaction

12 Research sessions: spatial data: Joining interval data in relational databases
Jost Enderle, Matthias Hampel, Thomas Seidl
June 2004 **Proceedings of the 2004 ACM SIGMOD international conference on Management of data**

Full text available: pdf(552.80 KB)   Additional Information: full citation, abstract, references

The increasing use of temporal and spatial data in present-day relational systems necessitates an efficient support of joins on interval-valued attributes. Standard join algorithms do not support those data types adequately, whereas special approaches for interval joins usually require an augmentation of the internal access methods which is not supported by existing relational systems. To overcome these problems we introduce new join algorithms for interval data. Based on the Relational Interval ...

13 Information systems interoperability: What lies beneath?
Jinsoo Park, Sudha Ram
October 2004 **ACM Transactions on Information Systems (TOIS)**, Volume 22 Issue 4

Full text available: pdf(824.78 KB)   Additional Information: full citation, abstract, references, index terms

Interoperability is the most critical issue facing businesses that need to access information from multiple information systems. Our objective in this research is to develop a comprehensive framework and methodology to facilitate semantic interoperability among distributed and heterogeneous information systems. A comprehensive framework for managing various semantic conflicts is proposed. Our proposed framework provides a unified view of the underlying representational and reasoning formalism ...

Keywords: Information integration, mediators, ontology, semantic conflict resolution, semantic heterogeneity

14 Special issue on spatial database systems: Management of multidimensional discrete data
Peter Baumann
October 1994 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 3 Issue 4

Full text available: pdf(2.30 MB)   Additional Information: full citation, abstract, references, citings

Spatial database management involves two main categories of data: vector and raster data. The former has received a lot of in-depth investigation; the latter still lacks a sound framework. Current DBMSs either regard raster data as pure byte sequences where the DBMS has no knowledge about the underlying semantics, or they do not complement array structures with storage mechanisms suitable for huge arrays, or they are designed as specialized systems with sophisticated imaging functionality, but n ...

Keywords: Multimedia database systems, image database systems, spatial index, tiling

15 Computing curricula 2001
September 2001 **Journal on Educational Resources in Computing (JERIC)**
Full text available: pdf(613.63 KB) html(2.78 KB)   Additional Information: full citation, references, citings, index terms

16 Advanced data processing in KRISYS: modeling concepts, implementation techniques, and client/server issues
Stefan Deßloch, Theo Härder, Nelson Mattos, Bernhard Mitschang, Joachim Thomas
May 1998   **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 7 Issue 2
Full text available: pdf(210.27 KB)   Additional Information: full citation, abstract, index terms

The increasing power of modern computers is steadily opening up new application domains for advanced data processing such as engineering and knowledge-based applications. To meet their requirements, concepts for advanced data management have been investigated during the last decade, especially in the field of object orientation. Over the last couple of years, the database group at the University of Kaiserslautern has been developing such an advanced database system, the KRISYS prototype. In this ...

Keywords: Client/server architectures, Consistency control, Object-oriented modeling concepts, Query processing, Run-time optimization

17 A fine-grained access control system for XML documents
Ernesto Damiani, Sabrina De Capitani di Vimercati, Stefano Paraboschi, Pierangela Samarati
May 2002   **ACM Transactions on Information and System Security (TISSEC)**, Volume 5 Issue 2
Full text available: pdf(330.60 KB)   Additional Information: full citation, abstract, references, citings, index terms

Web-based applications greatly increase information availability and ease of access, which is optimal for public information. The distribution and sharing of information via the Web that must be accessed in a selective way, such as electronic commerce transactions, require the definition and enforcement of security controls, ensuring that information will be accessible only to authorized entities. Different approaches have been proposed that address the problem of protecting information in a Web ...

Keywords: Access control, World Wide Web, XML documents, authorizations specification and enforcement

18 Query processing techniques in the summary-table-by-example database query language
Gultekin Özsoyoğlu, Victor Matos, Meral Özsoyoğlu
December 1989 **ACM Transactions on Database Systems (TODS)**, Volume 14 Issue 4
Full text available: pdf(3.52 MB)   Additional Information: full citation, abstract, references, citings, index terms, review

Summary-Table-by-Example (STBE) is a graphical language suitable for statistical database applications. STBE queries have a hierarchical subquery structure and manipulate summary tables and relations with set-valued attributes. The hierarchical arrangement of STBE queries naturally implies a tuple-by-tuple subquery evaluation strategy (similar to the nested loops join implementation technique) which may not be the best query processing strategy. In this paper we discuss the query ...

19 Ontology-based metadata: transforming the MARC legacy
Peter C. Weinstein
May 1998 **Proceedings of the third ACM conference on Digital libraries**
Full text available: pdf(1.21 MB)   Additional Information: full citation, references, citings, index terms

20 TIMBER: A native XML database
H. V. Jagadish, S. Al-Khalifa, A. Chapman, L. V. S. Lakshmanan, A. Nierman, S. Paparizos, J. M. Patel, D. Srivastava, N. Wiwatwattana, Y. Wu, C. Yu
December 2002 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 11 Issue 4
Full text available: pdf(268.39 KB)   Additional Information: full citation, abstract, index terms

This paper describes the overall design and architecture of the Timber XML database system currently being implemented at the University of Michigan. The system is based upon a bulk algebra for manipulating trees, and natively stores XML. New access methods have been developed to evaluate queries in the XML context, and new cost estimation and query optimization techniques have also been developed. We present performance numbers to support some of our design decisions. We believe that the key in ...

Keywords: Algebra, Document management, Hierarchical, Query processing, Semi-structured

Results 1 - 20 of 200   Result page: **1** 2 3 4 5 6 7 8 9 10   next

Useful downloads: Adobe Acrobat   QuickTime   Windows Media Player   Real Player